



Graphics card manufacturers boast about their products' 3D capabilities, but how close are we to real three-dimensional computer technology?
Mike Bedford investigates

The Victorians invented 3D photography over 100 years ago. Horror movies in 3D were all the rage in the 1950s, and the *Magic Eye* books of the 1990s brought 3D to a new generation. In recent years, computers have embraced 3D too. Architects use head-mounted displays to walk through buildings that are still on the drawing board and chemists can view molecular images rotating onscreen in 3D.

These technologies are expensive, the equipment isn't particularly user-friendly and it's aimed more at the engineering workstation than the PC. Times are changing, though. Prices are starting to fall and applications from games to graphic art are being enabled for 3D. The third dimension is coming to a PC screen near you.

STEREO VISION

For years graphics card manufacturers have claimed that their cards have 3D capability, but this just means that the cards can produce photo-realistic images by simulating perspective and shading. In this feature we look at real 3D, where objects seem to pop out of the screen.

The human visual system uses various techniques to gauge distance. A vital element is the ability to see a scene from two slightly different viewpoints. This is possible because we have two

forward-facing eyes. In a computer system, these two views have to be generated independently and the viewing hardware has to make sure that each eye sees just one of those two images.

ALTERNATING FRAMES

Most professional 3D systems display the left-eye and right-eye images alternately in rapid succession. Special glasses worn by the viewer then block alternate images so that the left eye sees only the left image and the right eye sees only the right image. This is called a frame sequential or field sequential display. There are two ways of doing it.

In one method the special glasses have LCD shutters instead of lenses. These shutters are similar to the LCD panels used in laptops and monitors, and they can be transparent or opaque, depending on the electrical signal sent to them. When the left image is displayed onscreen, the electronics in the glasses cause the left lens to become transparent and the right lens opaque. When the right image is displayed, the configuration is reversed. Each eye sees only one image and the brain perceives a three-dimensional effect. These glasses are known as active because they are driven by the PC. In the early days they were connected using a lead, but most modern active glasses are battery-operated and communicate with the PC via a wireless link.



Active glasses are heavy, and a method that uses passive glasses is now more popular. A special LCD panel is attached to the front of the monitor, which polarises the light using a circulatory polarising filter. Although it is technically more complex than plane polarisation, the effect is the same. The filter alternately polarises light in clockwise and anticlockwise directions in sync with the monitor's output for left and right eyes. The user wears glasses that polarise the light to each eye in different directions, filtering out the light intended for the other eye, so he sees a different image in each eye.

There are two requirements for graphics cards that are used as part of a frame sequential system. First they must provide what is known as quad buffered stereo so that the left and right images can be stored on the card. Only this way can they be swapped onscreen quickly enough to avoid flicker. Second, in order to synchronise the active glasses or the polarising panel with the alternating onscreen images, the cards must have a stereo connector to provide a signal to trigger the additional hardware. Many professional graphics cards have these facilities, but they are far from universal.

Frame sequential systems work with CRT monitors but not flat-screen displays. This is because the LCD panels in a flat-screen monitor are too slow for the swapping of the left- and right-eye images. If you used an LCD monitor you'd end up with ghosting because the left image would still be partially visible when the right image was supposed to be displayed, and vice versa.

Frame sequential systems are the sole domain of professionals with serious 3D applications. Prices put them out of reach of most other PC users. If you want to enter the world of 3D but have a modest budget, you might find lenticular systems more suitable.

LENTICULAR DISPLAYS

The 3D postcards that you can view without special glasses are usually lenticular cards. The left and right eye images are split into thin vertical stripes and alternated side by side. A sheet of thin, half-cylinder shaped lenses called lenticules attached to the front of the card directs the light from alternate stripes to your left and right eyes. Lenticular displays often include half a dozen or more images so you can see the scene from different viewpoints as you move your head.

The company 3DZ (www.3dz.co.uk) sells 3D lenticular kits for LCD monitors from just £152 including VAT. We asked company founder Phil Burridge what advantages it offers over field sequential displays. "This is the next stage in the evolution of display technology", he said. "It allows the user to view 3D content without wearing 3D glasses. But lenticular technology consumes much more computing power [to generate all those images] than a field sequential display and the resolution is lower; this is why most professionals are still using the field sequential-based displays."

Burridge explained how, because its technology splits up the images into diagonal stripes and uses special diagonal lenticules, the reduction in resolution is shared between the horizontal and vertical directions. The effective resolution of a 1,280x1,024 display when displaying a nine-image lenticular display is just 427x341 pixels, but the horizontal resolution of a conventional system would reduce it to 114 pixels.

HOW DO WE SEE IN 3D?

Take a look at the image of the three balloons below and try to decide which is the closest and which is the furthest away. If you think that the red balloon is the closest and the blue one the most distant, you'd be wrong, as you can see by looking at the plan view (right). The problem arises if you assume that all the balloons are the same size, so you use your understanding of perspective to calculate their relative distances. This shows how perspective alone can be misleading. Fortunately, our eyes use a number of other techniques to judge distances.



The key to our 3D vision is that we have two eyes that observe things from slightly different viewpoints, an effect that scientists call binocular disparity. This allows us to measure the angle from each eye to any object in a scene. Since our eyes are a fixed distance apart, this means that distances can be calculated by triangulation.

None of this is conscious, of course, but the end result is a perception of depth. Look at your finger and slowly touch your nose – you'll find yourself going cross-eyed in order to see it with both eyes without ending up with double vision. If you look at a distant object your eyes can see it clearly just by looking straight ahead.

All the 3D technologies discussed in this feature use binocular disparity, but some methods provide more visual depth clues. Try looking at a scene and moving your head from

side to side. You'll notice that nearby objects sway from side to side, whereas the background stays in the same place. Some objects that were previously obscured by a nearer object appear and others become eclipsed, and you'll be able to see around the sides of nearby objects.

Holographic images are able to reproduce this effect, providing a different image not just for two eyes but for multiple points, allowing the viewer to walk around the hologram and see a continually changing 3D image.



"With the computing power of CPU and 3D graphic chips dramatically advancing all the time, together with the latest ultra-resolution flat-panel technology, we strongly believe the age of auto-stereoscopic display [for which you don't have to wear special glasses] is coming," said Burridge.


The 3DZ kit is bundled with two pieces of software. The first can turn any 2D image into 3D. In its Manual Mode the user identifies objects in the scene, the software traces around them and the user assigns distances to those objects. The end result is a layered 3D effect. The

software's Automatic Mode attempts to do all this autonomously by analysing the scene. Although the end result is potentially more realistic than the layered effect, it isn't foolproof.

The other software is a driver that allows DirectX games to work in full 3D. This is available only for nVidia GeForce cards at the moment, but drivers for other cards are under development.

PARALLAX BARRIERS

Sharp's Actius AL3D notebook and its LL1513D monitor produce 3D images using a method that



is similar to 3DZ's lenticular kit. However, the two techniques are not identical. Sharp's technology caters only for stereo pairs – not the multiple images of lenticular displays – but this means each complete image has a higher resolution. There are other advantages, too, as Adrian Jacobs, research supervisor at Sharp Laboratories in Oxford, explained. “The Sharp 3D system uses a precision optical component called a parallax barrier. This barrier allows the left eye to see one column of pixels and the right eye to see another. A lenticular display uses a lens array fixed to the front of the display to create the same effect.”

What are the advantages of a parallax barrier over a lenticular display? According to Jacobs, “Sharp Laboratories has developed a parallax barrier that can be switched on and off electronically. This patented development means that a display can show full-resolution 2D for reading email or word processing and be switched to 3D for applications such as design, games or graphics.

“In effect the display has two LCD panels. The panel at the front of the screen is a standard TFT LCD. The panel at the rear has no pixels on it and doesn't display any information. However, when this panel is switched on it produces a latent parallax barrier: a series of vertical stripes that are alternately opaque and transparent. This is the 3D mode. When the rear panel is switched off, the parallax barrier is no longer formed and a 2D display results. As the barrier is behind the TFT LCD, the 2D image quality is perfect.”

SOLID PRINTING

Most 3D techniques fool the brain into seeing depth, but there is a computer-controlled device that creates a genuine three-dimensional solid object. These devices, originally used in industry with CAD software to prototype parts, are called rapid prototypers or solid printers.

As with two-dimensional printers there are a few different techniques, each with advantages and disadvantages. The first type, and one of the most common, uses a technique called stereo-lithography. This involves laser technology, but

that's where the similarity with office laser printers ends. Objects are built up a layer at a time, so the first job – which is carried out by software – is to split the computer model of the object into slices a fraction of a millimetre thick. As the shape of each slice is calculated, it's passed to the printer, which comprises a movable platform in a tank of liquid photopolymer resin and a moving ultraviolet laser.

Initially, the platform is just below the surface of the liquid. The laser scans the top surface of the resin to define the shape of the bottom slice. Wherever the laser beam illuminates the resin, it solidifies because of a process called cross-linking, in which additional chemical bonds are formed between the molecular strands in the liquid. The chemical process is similar to that which causes hardening in an epoxy adhesive once the hardener is added.

Since the platform was initially just below the surface of the photopolymer resin, this first slice will now be resting on the platform, which is lowered by the thickness of the slice to take it below the surface. Each subsequent slice is formed in the same way, adhering to the previous slice and building up a solid object in the process.

Overhangs are a potential problem. Imagine creating a solid model of your hand with your fingers pointing down and slightly apart. The first part of the model to be printed will be the tip of your middle finger, which will be resting on the platform. When the printer starts on the tip of your index finger it will appear as an island, unconnected to the rest of the layer and with nothing below it. As soon as it solidifies, therefore, it will drop down to the bottom of the tank. To avoid this, the software that splits the solid model into slices will look for any islands and add struts to keep the whole thing connected. When the model is finished all the struts are removed.

3D printers differ in the size of the output, their speed and their resolution. A developer of stereo-lithography and one of the major players in this market, 3D Systems (www.3dsystems.com), produces desktop 3D printers that can create objects about 200mm³ in size through to large floor-standing units that can work up to a

half-metre cube. The resolution will typically be 600dpi in each dimension, although the top-of-the-range SLA 7000 creates slices just 0.0254mm thick to achieve a 1,000dpi vertical resolution.

If you want to avail yourself of this technology, patience and a lot of money are essential. Entry-level desktop 3D printers, which use a lamination method rather than the stereo-lithography described here, cost \$23,000 (around £12,700), and you could pay \$700,000 (around £386,000) for a top-of-the-range stereo-lithographic unit. Although you might be able to create a small object in a couple of hours, large, complicated objects take several days to produce.

THE HOLOGRAPHIC REVOLUTION

Most methods of presenting a 3D electronic image involve preparing two or more images, each representing a different viewpoint, and arranging for a pair of them to be viewed at any one time, one by the left eye, the other by the right. Holography is different.

Traditional holograms are produced by illuminating an object with a laser beam and



It might sound like science fiction, but 3D Systems' top-of-the-range SLA 7000 causes solid objects to materialise in a tank of liquid



Spatial Imaging's Light Gate printer generates holographic output using two laser beams to create optical interference patterns

allowing the reflected light to fall on a high-resolution photographic film. The film is also illuminated by the same laser beam directly from the laser. No lenses are used so a hologram isn't an image in the normal sense of the word; instead it comprises millions of light and dark lines that are produced as the two light sources interfere with each other to produce constructive and destructive optical interference. The pattern of interference fringes acts as an ultra-complicated lens.

When you illuminate the processed hologram with an ordinary laser beam at the same angle as the direct laser beam that was used to create it, the hologram will produce a facsimile of the laser beam that had been reflected off the real object. Once again, this isn't an ordinary image; it is the pattern of light reflecting off every point in the object in every direction. So each eye automatically

sees the pair of images necessary for 3D vision and if you move your head around you see the scene from different positions. You can even look around objects to see what's behind them.

The interference patterns in a hologram are defined by the laws of physics, so it's possible to calculate them mathematically from a series of 2D images, each representing a different viewpoint. Because the lines in a hologram are separated by dimensions comparable to the wavelength of light – typically 0.5 microns (thousandths of a millimetre) – and because a resolution of a tenth of this – 0.05 microns – is needed to be able to represent them at any angle without jagged edges, the processing and printing requirements are phenomenal. It could easily take 20 hours to calculate the interference patterns for a one-inch square hologram. Meanwhile, printing them on film requires a multi-million pound electron beam plotter.

PIXEL PERFECT

However, so-called pixel holograms make this much more affordable, as Rob Munday, CEO of Spatial Imaging (www.holograms.co.uk) and expert in holography, explained. Instead of calculating all the interference fringes individually, an approximation to a hologram is created using two separate laser beams to produce linear interference patterns on a pixel-by-pixel basis.

A hologram isn't an image so those pixels don't correspond to points in the scene but, nevertheless, the end result is a hologram that can display parallax in the horizontal direction only. In other words, you can see 3D and you can get different views of the scene by moving your head from side to side but not up and down.

Even this has traditionally been a time-consuming process. Because of the need to allow any physical vibrations in the printer to settle down to a fraction of a micron before printing each pixel, 20 pixels per second has been a typical speed. Spatial Imaging has improved this throughput by a

factor of 100, as Munday explained.

"By decreasing the exposure time by a factor of 500, from 25 milliseconds to 50 microseconds, we can move the recording medium continuously rather than in a stop-start motion. The key, however, was how to move the two laser beams to different positions 2,000 times per second while keeping them in perfect registration with each other." Even with this technology, though, holographic printing isn't exactly speedy. At a resolution of 1,000dpi, an A4 hologram would take over 13 hours to print.

So does Munday see the holographic printer becoming as common as the laser printer is today? "I'm sure holography will eventually migrate to the desktop. However, unlike photography, it may never be developed as a mainstream hard-copy medium first. It may jump directly to an electronic medium like an LCD, but one that can resolve high-resolution interference fringes."

LOOKING FORWARD

Because we are creatures with two eyes that are used to seeing the world in 3D, it's surely a failure of technology that our PCs have been stuck in two dimensions for so long. Changes are afoot, but sequential frame and lenticular display systems are still a long way from the holographic display units much loved of science fiction film-makers.

Will moving, computer-generated holographic displays with horizontal and vertical parallax ever break out of the research labs? Most pundits are reluctant to say never. Still, it's sometimes good to take a reality check. As we've described above, it takes 20 hours of processing to calculate the interference fringes for a one-inch square hologram. Bearing in mind that a 17" display has an area of 130 square inches and that a moving display needs refreshing 50 times each second, it would seem that the technology is unlikely to be available on our desktops before a 1.4 Exahertz (or one billion gigahertz) processor. ☐

DIY 3D

If you're intrigued by the idea of three-dimensional images on your PC but don't want to spend a fortune, you can create your own 3D photos and view them onscreen with no special equipment whatsoever.

Taking the photographs is easy. Just take one photograph, move the camera 70mm to the right and take another shot. Make sure nothing in the photograph moves between shots. To help you keep the camera at the same height, you could try making a bracket to slide your camera along and mount it on a tripod. Now all you need is a means of viewing this stereo pair in 3D.

If you can do the ocular gymnastics (if you can view Magic Eye images you'll probably manage it), the easiest way is to display the two images side by side on your PC screen. You can do this with any decent image-editing software. This is similar to the technique used in Victorian stereograms, but since you're going to view them unaided, you'll have to restrict the size of each photograph to 70mm wide,

which is about 250 pixels on a 17" TFT screen. View them at a comfortable distance and allow yourself to go cross-eyed until the two images are exactly superimposed. At this point you'll actually see three images – the superimposed one in the centre plus one at either side – but they'll be blurred. Now – and this is the hard part – concentrate on the centre image and try to relax your eyes until it pops into focus but without losing convergence. If you manage this visual feat the centre image will appear in full 3D.

The second method is to create an anaglyph. This consists of one photo in shades of red and another in shades of cyan displayed on top of each other. When viewed through a pair of red-cyan 3D glasses, the red lens makes the white background appear red so the red image effectively disappears, leaving the cyan image appearing dark. Meanwhile, the cyan lens makes the cyan image disappear and the red image appear dark.

You can create an anaglyph from a stereo pair of photographs using photo-editing software, but you'll find it easier to use a dedicated program. Z-Anaglyph, which you can download free from www.z-graphix.com/anaglyph/zanag_en.htm,

is recommended. You'll find details in our guide 'How To... Create 3D Images' (*Shopper* September 2005). You can buy two pairs of red-cyan glasses for £1.50 including VAT from Media 55 (www.media55.co.uk/5.htm). Compared to side-by-side viewing, the major advantages of the anaglyph are that the image can be as large as you want and, since there's no knack involved, almost anyone can view them. The downside is that the colours can look a bit odd.



Create your own 3D images with Z-Anaglyph